IN THE CLAIMS:

Please cancel claim 27 and amend the claims 1-4, 6, 8, 10, 12, 13, 15, 16, 17 and 32 as follows.

1. (Currently Amended) An apparatus comprising:

a plurality of users;

resources that are partitioned according to a ranking of bandwidth associated with users, wherein the resources are partitioned according to a highest bandwidth supported by a node and an amount of bandwidth provided to each of the plurality of users is ranked from highest to lowest; and

a queue scheduler that

a) schedules one or more packets within the node during scheduling cycles, wherein each scheduling cycle is partitioned into regions that are coextensive with the highest bandwidth supported by the node and each schedule cycle is coextensive with a highest counting modulo partitions, and

b) services users associated with the highest bandwidth in at least one partition during each scheduling cycle and services consecutive bandwidth partitions of user associated with lower bandwidths across several cycles, wherein a number of scheduling cycles between servicing of consecutive bandwidth partitions increases as the bandwidth associated with the user decreases and the partition spacing for servicing a lower bandwidth user is determined by multiplying a number of lower bandwidth users that can

be serviced by the next highest bandwidth by a partition modulo of the next highest bandwidth.

~~a) distributes a partition worth of bandwidth to a plurality of queues according to a weight assigned to each of said queues so that each of said queues has its own sub-partition worth of data, each of said queues being capable of holding one or more packet identifiers, said plurality of queues arranged from a highest priority to a lowest priority, said queues serviced by a scheduler until each of said corresponding weights is consumed for each queue, and wherein higher priority queues are serviced before lower priority queues;~~

~~b) controls a flow of one or more packet identifiers from an active populated queue, until either~~

~~1) its unpopulated if less than its sub-partition worth of data has flowed,~~

~~2) its sub-partition worth of data has flowed, or~~

~~3) the combination of flows from those of said queues that have been active results in said partition worth of data having flowed from said those of said queues that have been active, as a whole,~~

~~wherein a populated queue is deemed active if it is the highest priority populated queue out of those of said populated queues that have not yet been deemed active, such that populated queues are deemed active in succession until the lowest priority populated queue has been deemed active or until the combination of flows from those of~~

~~said queues that have been active results in said partition worth of data having flowed from those of said queues that have been active, as a whole.~~

2. (Currently Amended) A method comprising:

a) . dividing a total amount of data, based upon an individual weight assigned to each of a plurality of queues, into an amount of data that each of said queues may service;

b) scheduling packets in the plurality of queues during scheduling cycles, wherein each scheduling cycle is partitioned into regions that are coextensive with a highest bandwidth being managed by a node and each schedule cycle is coextensive with a highest counting modulo partitions;

c) servicing queues associated with the highest bandwidth in at least one partition during each scheduling cycle and servicing consecutive bandwidth partitions of queues associated with lower bandwidths across several cycles, wherein a number of scheduling cycles between servicing of consecutive bandwidth partitions increases as the bandwidth associated with the queue decreases and partition spacing for servicing a lower bandwidth queue is determined by multiplying a number of lower bandwidth users that can be serviced by the next highest bandwidth by a partition modulo of the next highest bandwidth;

d) servicing one or more populated queues, each of said servicing of a populated queue continuing until said populated queue is no longer populated or said amount of data determined for said populated queue has been released; and

e) servicing one or more of said queues that remain populated, if any, until said total amount of data has been released from all of said queues in combination including said servicing of said populated queues.

3. (Currently Amended) A method, comprising

a) distributing a partition worth of data across a plurality of queues according to a weight assigned to each of said queues so that each of said queues has its own sub-partition worth of data, each of said queues capable of holding one or more packet identifiers, each of said one or more packet identifiers pointing to its own packet, said plurality of queues ranging from a highest priority queue to a lowest priority queue; and

b) scheduling packets in the plurality of queues during scheduling cycles, wherein each scheduling cycle is partitioned into regions that are coextensive with a highest bandwidth being managed by a node and each schedule cycle is coextensive with a highest counting modulo partitions, and

c) servicing queues associated with the highest bandwidth in at least one partition during each scheduling cycle and servicing consecutive bandwidth partitions of queues associated with lower bandwidths across several cycles, wherein a number of scheduling cycles between servicing of consecutive bandwidth partitions increases as the bandwidth

associated with the queue decreases and partition spacing for servicing a lower bandwidth queue is determined by multiplying a number of lower bandwidth users that can be serviced by the next highest bandwidth by a partition modulo of the next highest bandwidth; and

d) flowing a flow of one or more packet identifiers from an active populated queue, until:

1) its unpopulated if less than its sub-partition worth of data has flowed, or until

2) its sub-partition worth of data has flowed, or until

3) the combination of flows from those of said queues that have been active results in said partition worth of data having flowed from said those of said queues that have been active, as a whole.

~~wherein a populated queue is deemed active if it is the highest priority populated queue out of those of said populated queues that have not yet been deemed active, such that populated queues are deemed active in succession until the lowest priority populated queue has been deemed active or until the combination of flows from those of said queues that have been active results in said partition worth of data having flowed from those of said queues that have been active, as a whole.~~

4. (Currently Amended) The method of claim ~~1~~3 wherein each of said packet identifiers further comprise the same Port_ID value that identifies a port to which said queues belong.

5. (Previously Presented) The method of claim 4, wherein said port handles packets destined to the same user.

6. (Currently Amended) The method of claim ~~1~~3 wherein one of said queues receives only those of said packet identifiers that each point to its own networking control packet.

7. (Previously Presently) The method of claim 4 wherein said one of said queues is said highest priority queue.

8. (Currently Amended) The method of claim ~~1~~3 wherein one of said queues receives only those of said packet identifiers that each point to its own real time traffic packet.

9. (Previously Presently) The method of claim 6 wherein said one of said queues is a second highest priority queue.

10. (Currently Amended) The method of claim ~~1~~3 wherein one of said queues receives only those of said packet identifiers that each point to its own fast data traffic packet.

11. (Previously Presently) The method of claim 8 wherein said one of said queues receives only those of said packet identifiers that each point to its own traditional data traffic packet.

12. (Currently Amended) The method of claim ~~1~~3 wherein said partition worth of data is a scheduling cycle partition worth of data, wherein one scheduling cycle partition worth of data per scheduling cycle corresponds to a data rate that is a highest data rate managed by a networking system to which each of said queues belong.

13. (Currently Amended) The method of claim ~~1~~3 wherein each of said weights add up to a value that represents 100% or less of said partition worth of data.

14. (Previously Presently) The method of claim 11 wherein each of said weights are equal.

15. (Currently Amended) The method of claim ~~1~~3 wherein each of said weights add up to a value that represents more than 100% of said partition worth of data.

16. (Currently Amended) The method of claim ~~13~~, wherein if more than an active queue's sub-partition worth of data had flowed while it was active, the difference between the amount of data that flowed and said sub-partition worth of data is subtracted from said active queue's sub partition worth of data in order to reduce the flow the next time said active queue becomes active.

17. (Currently Amended) A method, comprising

a) distributing a partition worth of data across a plurality of queues according to a weight assigned to each of said queues so that each of said queues has its own sub-partition worth of data, each of said queues capable of holding one or more packet identifiers, each of said one or more packet identifiers pointing to its own packet, said plurality of queues ranging from a highest priority queue to a lowest priority queue;

b) scheduling packets in the plurality of queues during scheduling cycles, wherein each scheduling cycle is partitioned into regions that are coextensive with a highest bandwidth being managed by a node and each schedule cycle is coextensive with a highest counting modulo partitions, and

c) servicing queues associated with the highest bandwidth in at least one partition during each scheduling cycle and servicing consecutive bandwidth partitions of queues associated with lower bandwidths across several cycles, wherein a number of scheduling cycles between servicing of consecutive bandwidth partitions increases as the bandwidth associated with the queue decreases and partition spacing for servicing a lower bandwidth

queue is determined by multiplying a number of lower bandwidth users that can be serviced by the next highest bandwidth by a partition modulo of the next highest bandwidth;

d) flowing a flow of one or more packet identifiers from an active populated queue, until:

1) its unpopulated if less than its sub-partition worth of data has flowed, or until

2) its sub-partition worth of data has flowed, or until

3) the combination of flows from those of said queues that have been active results in said partition worth of data having flowed from said those of said queues that have been active, as a whole,

wherein a populated queue is deemed active if it is the highest priority populated queue out of those of said populated queues that have not yet been deemed active, such that populated queues are deemed active in succession until the lowest priority populated queue has been deemed active or until the combination of flows from those of said queues that have been active results in said partition worth of data having flowed from said those of said queues that have been active, as a whole.~~; and~~

~~e) if:~~

~~1)one or more populated queues exist after each of said populated queues has been active, and~~

~~2) said combination of flows from those of said queues that have been active results in less than said partition worth of data having flowed from said those of said queues that have been active, as a whole;~~

~~then:~~

~~1) flowing one or more additional flows from said one or more populated queues until said partition worth of data has flowed from said queues as a whole, or until~~

~~2) each of said queues is unpopulated if each of said queues becomes unpopulated before said partition worth of data has flowed from said queues as a whole.~~

18. (Previously Presently) The method of claim 15 wherein each of said packet identifiers further comprise the same Port_ID value that identifies a port to which said queues belong.

19. (Previously Presently) The method of claim 16 wherein said port handles packets destined to the same user.

20. (Previously Presently) The method of claim 15 wherein one of said queues receives only those of said packet identifiers that each point to its own networking control packet.

21. (Previously Presently) The method of claim 18 wherein said one of said

queues is said highest priority queue.

22. (Previously Presently) The method of claim 15 wherein one of said queues receives only those of said packet identifiers that each point to its own real time traffic packet.

23. (Previously Presently) The method of claim 20 wherein said one of said queues is a second highest priority queue.

24. (Previously Presently) The method of claim 15 wherein one of said queues receives only those of said packet identifiers that each point to its own fast data traffic packet.

25. (Previously Presently) The method of claim 22 wherein said one of said queues receives only those of said packet identifiers that each point to its own traditional data traffic packet.

26. (Previously Presently) The method of claim 15 wherein said partition worth of data is a scheduling cycle partition worth of data, wherein one scheduling cycle partition worth of data per scheduling cycle corresponds to a data rate that is a highest data rate managed by a networking system to which each of said queues belong.

27. (Cancelled)

28. (Previously Presently) The method of claim 25 wherein each of said weights are equal.

29. (Cancelled)

30. (Previously Presently) The method of claim 15 wherein said flowing one or more additional flows further comprises flowing packet identifiers from a next queue, said next queue following a previous queue that flowed an additional flow to consume a previous distribution of a partition worth of data.

31. (Previously Presently) The method of claim 15 wherein if more than an active queue's sub-partition worth of data had flowed while it was active, the difference between the amount of data that flowed and said sub-partition worth of data is subtracted from said active queue's sub partition worth of data in order to reduce the flow the next time said active queue becomes active.

32. (Currently Amended) An apparatus, comprising:

a scheduler that

~~a) distributes a partition worth of data across a plurality of queues according to a weight assigned to each of said queues so that each of said queues has its own sub-partition worth of data, each of said queues capable of holding one or more packet identifiers, each of said one or more packet identifiers pointing to its own packet, said plurality of queues ranging from a highest priority queue to a lowest priority queue; and~~

b) scheduling packets in the plurality of queues during scheduling cycles, wherein each scheduling cycle is partitioned into regions that are coextensive with a highest bandwidth being managed by a node and each schedule cycle is coextensive with a highest counting modulo partitions, and

c) servicing queues associated with the highest bandwidth in at least one partition during each scheduling cycle and servicing consecutive bandwidth partitions of queues associated with lower bandwidths across several cycles, wherein a number of scheduling cycles between servicing of consecutive bandwidth partitions increases as the bandwidth associated with the queue decreases and partition spacing for servicing a lower bandwidth queue is determined by multiplying a number of lower bandwidth users that can be serviced by the next highest bandwidth by a partition modulo of the next highest bandwidth; and

b) controls a flow of one or more packet identifiers from an active populated queue, until:

1) its unpopulated if less than its sub-partition worth of data has flowed, or until

2) its sub-partition worth of data has flowed, or until

3) the combination of flows from those of said queues that have been active results in said partition worth of data having flowed from said those of said queues that have been active, as a whole,

wherein a populated queue is deemed active if it is the highest priority populated queue out of those of said populated queues that have not yet been deemed active, such that populated queues are deemed active in succession until the lowest priority populated queue has been deemed active or until the combination of flows from those of said queues that have been active results in said partition worth of data having flowed from said those of said queues that have been active, as a whole.

33. (Previously Presently) The apparatus of claim 30 wherein each of said packet identifiers further comprise the same Port_ID value that identifies a port to which said queues belong.

34. (Previously Presently) The apparatus of claim 31 wherein said port handles packets destined to the same user.

35. (Previously Presently) The apparatus of claim 30 wherein one of said queues receives only those of said packet identifiers that each point to its own networking control packet.

36. (Previously Presently) The apparatus of claim 33 wherein said one of said queues is said highest priority queue.

37. (Previously Presently) The apparatus of claim 30 wherein one of said queues receives only those of said packet identifiers that each point to its own real time traffic packet.

38. (Previously Presently) The apparatus of claim 35 wherein said one of said queues is a second highest priority queue.

39. (Previously Presently) The apparatus of claim 30 wherein one of said queues receives only those of said packet identifiers that each point to its own fast data traffic packet.

40. (Previously Presently) The apparatus of claim 37 wherein said one of said queues receives only those of said packet identifiers that each point to its own traditional data traffic packet.

41. (Previously Presently) The apparatus of claim 30 wherein said partition worth of data is a scheduling cycle partition worth of data, wherein one scheduling

cycle partition worth of data per scheduling cycle corresponds to a data rate that is a highest data rate managed by a networking system to which each of said queues belong.

42. (Previously Presently) The apparatus of claim 30 wherein each of said weights add up to a value that represents 100% or less of said partition worth of data.

43. (Previously Presently) The apparatus of claim 40 wherein each of said weights are equal.

44. (Cancelled)

45. (Previously Presently) The apparatus of claim 30 wherein if more than an active queue's sub-partition worth of data had flowed while it was active, said scheduler takes the difference between the amount of data that flowed and said sub-partition worth of data and subtracts it from said active queue's sub partition worth of data in order to reduce the flow the next time said active queue becomes active.